CIS-TR-81-5

A SURVEY OF GOSSIPING AND BROADCASTING IN COMMUNICATION NETWORKS

by

Sandra M. Hedetniemi*

Stephen T. Hedetniemi*

Department of Computer and Information Science
University of Oregon
Eugene, OR 97403

Abstract

This survey reviews more than 50 papers concerning the mathematical study of information dissemination processes. Emphasis is given to the study of two of these processes, called gossiping and broadcasting. In gossiping, each person in a group has a unique item of information and must transmit it to everyone else. Thus gossiping is an all-to-all process. In broadcasting, only one person has an item of information which must be transmitted to everyone. Broadcasting is thus a one-to-all process.

* Research supported in part by the National Science Foundation under Grant ENG-79-02960

1. Introduction

In 1949 and 1950, Leavitt [38] and Bavelas [2], respectively, studied the effectiveness of different communication patterns in helping small groups of people solve common tasks. A typical task studied was the following: each of five subjects is dealt five playing cards; the cards may not be passed around, but the subjects may communicate with one another according to a given communication pattern by writing messages; the task is considered finished when each subject selects one of their five cards so that the five cards selected comprise the highest ranking poker hand that can be made by selecting one card from each person. (We assume that each subject has a perfect knowledge of poker hand ratings.) For communication patterns such as those indicated in Figure 1, Leavitt and Bavelas considered such measures as the number of messages required, the time required, the average number of errors made and the ratings of morale and satisfaction in completing such tasks.

Figure 1

In 1954 Landau [37] studied task-oriented groups in which each individual, initially, has one piece of information which must be transmitted to all the others in order to complete a task. At every sending time, each individual sends all the information they have acquired to one other individual. A typical example of a task in this case is the following: "each individual is given a set of colored marbles, only one color being common to all the group. The members of the group must exchange messages about their own colors and what they have learned about the colors of the others, until finally everybody knows the common color. ... Messages are sent only after everyone has indicated readiness to transmit; the transmissions then take place simultaneously, each individual sending to just one of their possible recipients. Each member of group knows initially to whom they can send messages, but does not know to whom the others can send." [37] A major assumption that Landau made was that each individual picks a recipient at random from among those to whom they are permitted to send messages, with equal probability of picking any one of them. For a variety of communication patterns among three and four people, Landau determined the expected value of the time it would take to complete such a task.

A natural step in the progression of studies of communication processes occured some eighteen years later when Hajnal, Milner and Szemerédi(1972) [28] attributed to A. Boyd the following problem: "There are n ladies, and each one of them knows an item of scandal which is not known to any of the others. They communicate by telephone, and whenever two ladies make a call, they pass on to each other, as much scandal as they know at the time. How many calls are needed before all ladies know all the scandal?" [28]. This problem, which has become known as the Gossip Problem, or the Telephone Problem, has in turn been the source of more than three dozen research papers that have studied problems concerning the spread of information among a set of people, whether it be by telephone calls, conference calls, letters or even computer networks.

- Information Dissemination
    - Gossiping
        - Two-way calls
            - complete graphs
            - arbitrary graphs
        - One-way calls
        - Conference calls
        - Random calls
    - Broadcasting
        - Local broadcasting
            - Single message broadcasting
                - arbitrary graphs
                - minimal broadcast graphs
            - Multiple message broadcasting
        - Long distance broadcasting
    - Grids
        - Broadcasting
        - Gossiping
    - Reliability
    - Broadcasting in computer networks
    - New areas of study
        - Shouting
        - Minimum cost long distance broadcasting
        - Weighted gossiping
        - 1-way communication
        - Polling
        - Multi-destination transmittal

Figure 2. Areas of study concerning information dissemination

contained the following result.

(1) Let f(n) equal the minimum number of calls necessary to complete gossiping among n people. Then f(n) satisfies the following: $f(1) = 0$; $f(2) = 1$; $f(3) = 3$ and $f(n) = 2n-4$, for $n \geq 4$.

This result was proved by R. T. Bumby [6] and J. Spencer (unpublished) (cf. [28]), Hajnal, Milner and Szemerédi (1972) [28], Tijdeman (1971) [50] and Baker and Shostak (1972) [1]. Each of the published proofs, varying in approach and length, provide interesting perceptions on the nature of the Gossip Problem.

The proof by Hajnal, Milner and Szemerédi is a particularly ingenious one that makes use of an interchange rule in complete calling sequences. This rule considers conditions under which calls in a particular sequence can be interchanged in time yet still complete gossiping. Their proof also focuses on the positions in a calling sequence by which time an individual knows everything. Let us call such an individual an <u>expert</u>. The following two results can be used to prove the main result ((1) above).

(2) In any sequence of calls among n people there are no experts after n-2 calls.

(3) In any calling sequence among n people, after n+k-4 calls, there are at most k experts and if there are exactly k experts, then there is an equivalent calling sequence in which the last k calls are made exclusively between the experts.

This is suggested by the observation:

"Suppose an additonal gossip g joins a group of n persons. If g phones any one of the gossips first, the additional information known by g will be transmitted in f(n) calls among the n gossips. After this, one additional call is needed by g to obtain all the gossip." [3].

One of the first variations of the Gossip Problem to be studied restricted the communication patterns among the people, as had Leavitt [38], Bavelas [2] and Landau [37]. Harary and Schwenk (1974) [29] and Golumbic (1974) [27] considered the situation where an individual can call some but not all other people. The original Gossip Problem, expressed in terms of the underlying communication graph, where vertices represent people and edges represent the allowable communication lines, assumes that the communication graph is complete, i.e. a call can be made between any two people. Harary and Schwenk, in placing restrictions on the underlying communication graphs, obtained the following results.

(7) If the communication graph among a set of n people is a tree, then $f(n) = 2n-3$, for $n \geq 2$.

(8) If the communication graph among a set of n people is a unicyclic graph, whose cycle is not a 4-cycle, then $f(n) = 2n-3$, for $n \geq 2$.

(9) For any connected, communication graph with n vertices, $2n-4 \leq f(n) \leq 2n-3$, for $n \geq 4$.

(10) For any connected, communication graph with n vertices, which contains a 4-cycle, $f(n) = 2n-4$, for $n \geq 4$.

(i) <u>a-redundant</u> if neither A nor B know everything before the call;

(ii) <u>1-redundant</u> if either A or B, but not both, knows everything before the call; and

(iii) <u>2-redundant</u> if both A and B know everything before the call.

Cot was interested in the existence of a-redundant calling sequences which complete gossiping in the minimum number (2n-4) of calls. He found two cases for which this was possible, i.e. when n = 4 and n = 8. Cot also considered briefly the case of gossiping in which each call has associated with it a given cost, depending on such factors as the distance between the two vertices or the number of messages transmitted. In particular, one is interested in finding a calling sequence which minimizes the total cost of gossiping.

Another generalization of the Gossip Problem is obtained by assuming that information is transmitted by 'conference' or 'k-party' calls. The following result was first proved by Lebensold (1973) [39], then by Bermond (1976) [4] and still later by Kleitman and Shearer (1980) [32], who reduced Lebensold's 13-page proof to an elegant 3-page proof.

(14) Let f(n,k) equal the minimum number of k-party calls necessary to complete gossiping among n people (in a complete communication graph). Then

to any given individual. Landau proved the following result.

(16) $t_r(n) = [\log_{r+1} n]$.

Entringer and Slater (1976) [13] have considered the minimum amount of time necessary to complete gossiping in complete digraphs, under two different conditions. First, let $t_2(n,k)$ equal the minimum number of time units necessary to complete gossiping subject to the constraint that during each period of time each person can send all the information they know to each of at most k other people <u>and</u> each of at most k other people can send information to them. Second, let $t_1(n,k)$ equal the minimum number of time units necessary to complete gossiping subject to the constraint that during each time unit a person can either send all the information they know to at most k other people <u>or</u> they can receive information from at most k other people. No one can both send and receive information during a time unit. They obtained the following two results.

(17) For any complete communication graph with n vertices,

(i) $t_2(n,k) = [\log_{k+1} n]$

(ii) $[\log_{k+1} n] \leq t_1(n,k) \leq 2[\log_{k+1} n]$

Still another variant of the Gossip Problem is obtained by restricting the allowable calling sequences so that no one hears their own information (cf. the paper by Baker and Shostak [1]). That is, no one can call a person if the caller already knows the unique piece of information originally known only by the person called. West (1978) [56] has studied this problem by determining

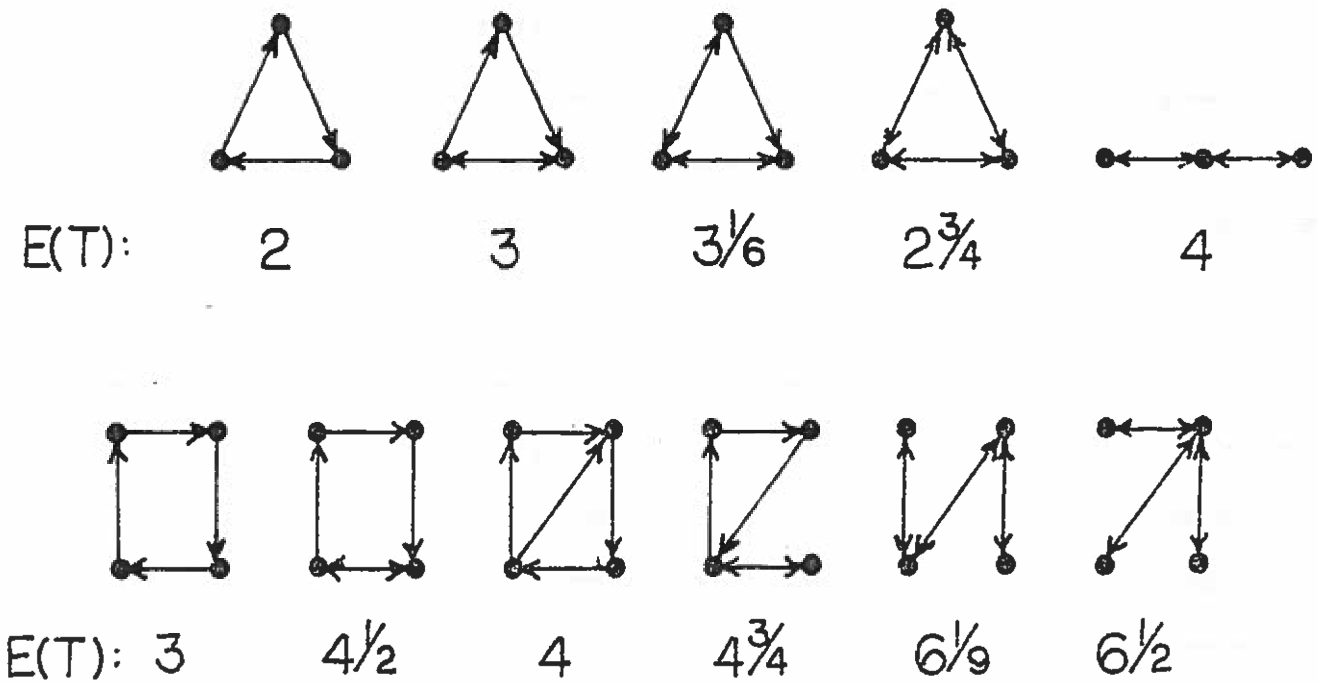

FIGURE 3. EXPECTED TIMES TO COMPLETE RANDOM 1-WAY GOSSIPING FOR VARIOUS COMMUNICATION GRAPHS

units required to complete broadcasting from vertex u?". We define the <u>broadcast time of vertex u</u>, b(u), to equal this minimum time. It is easy to see that for any vertex u in a connected graph G with n vertices, $b(u) \geq [\log_2 n]$, since during each time unit the number of informed vertices can at most double. We define the <u>broadcast time of a graph G</u>, b(G), to equal the maximum broadcast time of any vertex u in G, i.e. $b(G) = \max\{b(u) \mid u \in V(G)\}$. For the complete graph K_n with $n \geq 2$ vertices, $b(K_n) = [\log_2 n]$, yet K_n is not minimal with respect to this property. That is, we can remove edges from K_n and still have a graph G with n vertices such that $b(G) = [\log_2 n]$. We define a <u>minimal broadcast graph</u> to be a graph G with n vertices such that $b(G) = [\log_2 n]$, but for every proper spanning subgraph $G' \subset G$, $b(G') > [\log_2 n]$.

We define the <u>broadcast function</u> B(n) to equal the minimum number of edges in any minimal broadcast graph on n vertices. A <u>minimum broadcast graph</u> (mbg) is a minimal broadcast graph on n vertices having B(n) edges. From an applications perspective, minimum broadcast graphs represent the cheapest possible communication networks (having the fewest communication lines) in which broadcasting can be accomplished, from any vertex, as fast as theoretically possible.

We define a <u>minimum broadcast tree</u> to be a rooted tree with n vertices and root u such that $b(u) = [\log_2 n]$. In any connected graph G, a broadcast from a vertex u determines a spanning tree rooted at u. Thus, every vertex of a minimal broadcast graph is the root of a minimum broadcast tree which contains all of the vertices of G.

(21) For $1 \leq n \leq 15$, the values of the broadcast function $B(n)$ are as follows:

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
B(n):	0	1	2	4	5	6	7	12	10	12	13	15	18	21	24

Farley, Hedetniemi, Mitchell and Proskurowski [19] also found examples of minimum broadcast graphs for each value of n, $1 \leq n \leq 15$. These examples were extended by Mitchell and Hedetniemi (1980) [42] who found all mbgs with $n \leq 11$ vertices and presented the only known examples of mbgs with $n \leq 15$ vertices. Their census was later augmented by Bermond (1981) [5] who found a second mbg with 14 vertices. The number of mbgs is summarized in the following result.

(22) For $1 \leq n \leq 15$, the number of mbgs with n vertices is as follows:

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
mbgs:	1	1	1	1	4	1	1	6	6	21	2	1*	1*	2*	3*,

where an * indicates the number of mbgs known to exist.

The results of these studies suggest that mbgs are extremely difficult to find; in fact, no procedure is known for constructing an mbg with n vertices for any value of $n \geq 18$, except for the values $n = 2^k$, where mbgs are easy to construct. Figure 4 illustrates several examples of minimum broadcast graphs (with the author's names indicated in parentheses).

Figure 4

To further illustrate these mbgs, we present in Figure 5 Bermond's demonstration [5] that vertices 1 and 2 in the Heawood graph on 14 vertices in Figure 1 can broadcast to every other vertex in $\lceil \log_2 14 \rceil = 4$ time units, eg. at time 1 vertex 1 calls vertex 2 (or vice-versa), and at times 2 and 3 vertex 1 calls vertices 14 and 6, respectively.

Figure 5

In [15] Farley (1980) modified the local broadcasting restriction by permitting 'long distance' calls, i.e. a vertex u may call any vertex w if there exists a path from u to w, no edge of which is being used in any other call. He obtained a constructive proof of the following, rather nice result.

(23) In any connected graph G with n vertices, every vertex can complete long distance broadcasting in $\lceil \log_2 n \rceil$ time units.

Proskurowski (1981) [45] was the first to study minimum broadcast trees (mbts). In particular, he designed an O(n) algorithm for deciding if an arbitrary rooted tree with n vertices is and mbt; he designed an algorithm for generating all mbts with n vertices; and he developed a recurrence relation to count a number of subsets of mbts. By coincidence, all minimum broadcast trees with 2^k vertices are binomial trees B_k (cf. Vuillemin (1978) [53]), a class of trees which have the simple inductive definition indicated in Figure 6. Figure 6 also illustrates several binomial (or minimum broadcast) trees. (The tree in Figure 5, incidentally, is also a minimum broadcast tree on 14 vertices.)

Binomial trees occur naturally in the study of various data manipulation problems, in particular those involving priority queues (cf. [53]).

Figure 6

The next work on broadcasting reconsidered the assumption that each call requires one unit of time. In particular, if one is broadcasting large files of information over the lines of a computer network, then this assumption is not very realistic. It was quickly realized that if one needed to broadcast more than one message, then simple modifications of one-message broadcast schemes would be too inefficient. For example, neither repeating a one-message broadcast scheme k times, nor having each call take k time units produces a very efficient k-message broadcast scheme. It is easy, for example, to construct 'mixed' calling schemes which complete multiple-message broadcasting in less time than either of these two schemes (cf. Figure 7).

Figure 7

Consider the following functions:

P(m,t): the maximum number of people who can be informed of m messages in t time units;

M(p,t): the maximum number of messages which can be broadcast to p people in t time units; and

T(m,p): the minimum number of time units necessary to broadcast m messages to p people.

Chinn, Hedetniemi and Mitchell (1979) [7] and Farley [16]

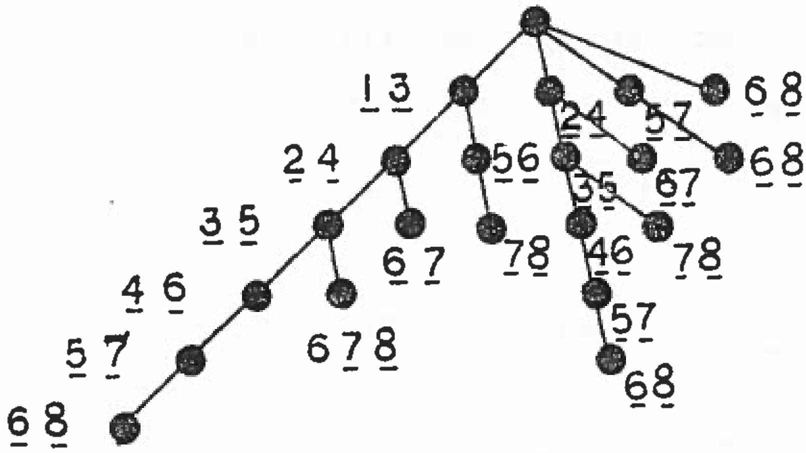

A. Efficient multiple message broadcasting

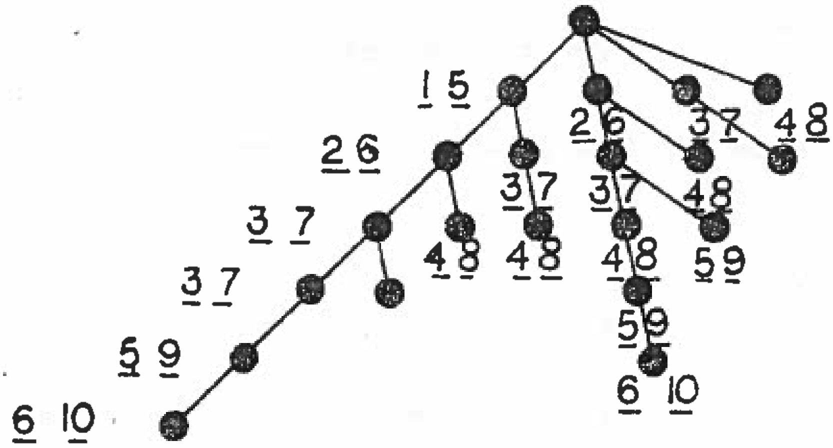

B. Repeating a single message broadcast twice

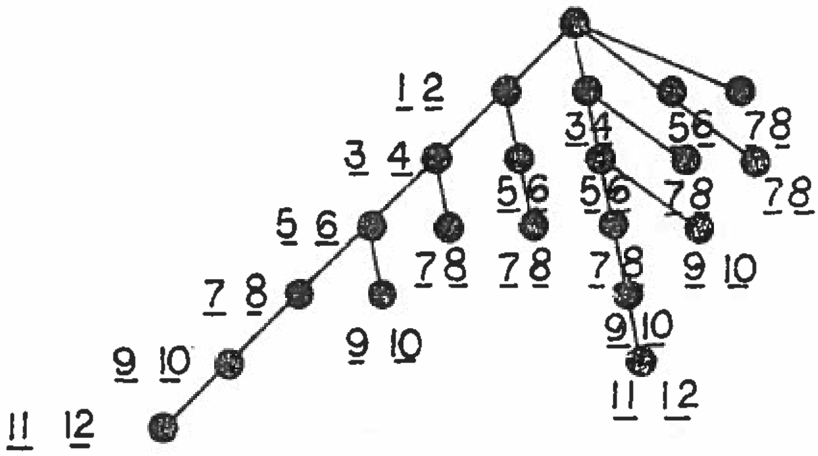

C. Sending both messages at the same time

Figure 7. Broadcasting Two Messages in a Tree by Three Different Methods

The next variant of local broadcasting to be studied considered the problem of determining the minimum number of message originators necessary to complete broadcasting in a specified amount of time. Hedetniemi and Hedetniemi (1979) [30] observed that if broadcasting must be completed in 1 time unit in an arbitrary communication graph, then the graph must be decomposed onto a minimum number of subgraphs, each of which is either a K_1 or a K_2 (this is equivalent to the Maximum Matching Problem). If broadcasting must be completed in 2 time units, then the graph must be decomposed into a minimum number of paths, each of which is of length at most 3.

In [30] they present a linear algorithm for decomposing an arbitrary tree into a minimum number of paths, each of which has length $\leq k$, for arbitary values of k. Shortly thereafter, Farley and Proskurowski (1980) [22] completely settled the question of determining the minimum number of originators necessary to complete broadcasting in an arbitrary tree in at most t time units, for arbitrary values of t. They present a linear algorithm for decomposing a tree into a minimum number of subtrees, in each of which broadcasting can be completed in at most t time units.

4. Grid graphs

Grid graphs (cf. Figure 8) are a special class of graphs which have received a lot of study from many different perspectives. They have been used, for example, to model games on a checkerboard, networks of city streets, telephone switching networks, geographical data bases, matrix manipulations, the cellular spaces of John von Neumann, and parallel computer architectures, eg. the ILLIAC IV

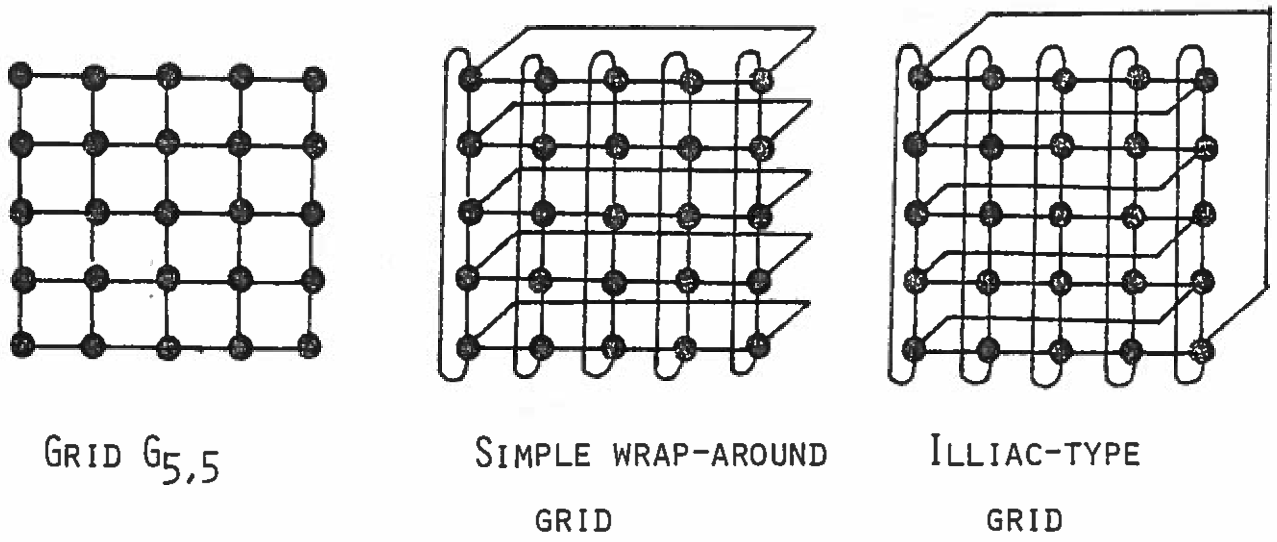

FIGURE 8. EXAMPLES OF GRID AND GRID-LIKE GRAPHS

their paper by conjecturing that the function $f(2,t)$, which equals the maximum number of vertices which can be informed of a single message originated by a given 'cell' in the infinite 2-dimensional grid, after t time units, by any local broadcasting scheme, is given by

$$f(2,t) = 2t^2 - 6t + 8, \quad \text{for } t \geq 2.$$

This conjecture gave rise to several subsequent papers. First, Cockayne and Hedetniemi (1978) [8] conjectured that in the n-dimensional grid,

$$f(n,t) = 2^t, \quad \text{when } t \leq 2n, \text{ and}$$

$$f(n, 2n+k+1) = 2^n \sum_{j=0}^{n} 2^j \binom{n-j+k}{k}, \quad \text{for other values of } t.$$

In his Ph.D. thesis (1979), Ko [34] and [35] later verified the 2-dimensional conjecture of Farley and Hedetniemi and obtained several other results as follows.

Let $f(n,t)$ represent the maximum number of cells which can be informed of a message originated by a given cell in an n-dimensional grid graph, after t time units, by a local broadcasting scheme.

(34) $f(2,t) = 2t^2 - 6t + 8, \quad \text{for } t \geq 2.$

(35) $f(n,t) \leq \frac{2^n}{n!} t^n + \frac{2 - n2^n}{(n-1)!} t^{n-1} + \frac{1}{(n-2)!} (4-n- \frac{1}{3} (2^{n+1}-n2^{n-2}) + n^2 2^{n-1}) t^{n-2} + O(t^{n-3}).$

(36) $f(3,t) = \frac{4}{3} t^3 - 11 t^2 + \frac{101}{3} t - 28, \quad \text{for } t \geq 9.$

grid graphs obtained by whispering (in which a cell can only call one neighbor per unit of time) and shouting (in which a cell can simultaneously call all of its neighbors). Let s(t) denote the number of cells informed by time t if information is spread from one cell at time t=0 by shouting, and let w(t) denote the maximum number of cells that can be informed by any whispering algorithm by time t. Stout proved the following result.

(42) Whispering is asymptotically as efficient as shouting in the sense that $\lim \; s(t)/w(t) = 1$.

In another paper on grid graphs, Van Scoy (1979) [52] considered broadcasting multiple mesages in complete grids $G_{m,n}$. In particular, she designed an algorithm for broadcasting k messages in $G_{m,n}$ in 2n + 2k - 4 units, where $k \leq n-2$ if n is odd and $k \leq n-1$ if n is even.

5. Reliability studies

Recently a few papers have been written concerning problems of broadcasting in the presence of faults, i.e. the problems which arise when one or more communication lines fail or one or more communication sites (vertices) are inoperative. Liestman (1980) [40], for example, has studied several parameters related to fault tolerant broadcasting. Let $T_{0,k}(n)$ equal the minimum time required to broadcast in the presence of k faults in any graph on n vertices. Liestman showed the following.

(43) $T_{0,1}(n) = \{\log_2 n\} + 1$.

(44) $T_{0,2}(n) = \{\log_2 n\} + 2$, for $n \geq 5$, $n \neq 4i + 3$.

6. Broadcasting in computer networks

Several papers have been published which are closely related to the study of gossiping and broadcasting, but which consider somewhat different models of information dissemination. Perhaps the most notable of these is the work of Dalal (1977) [11], Dalal and Metcalfe (1978) [12], Wall and Owicki (1980) [55] and Wall (1980) [54], who have considered special problems of broadcasting in packet-switched computer networks, like ARPANET. A fundamental problem with networks like ARPANET is that they are not built with a mechanism to handle broadcasting. Consequently, in order to carry out broadcasting from a given vertex an individual message must be routed to each of the other vertices in the network.

Dalal [11] and Dalal and Metcalfe [12] proposed several different broadcasting algorithms for overcoming this deficiency. A basic assumption in these algorithms is that once a vertex receives a message along a given communication line, it can, in effect, 'shout' that message (using Stout's terminology [49]), simultaneously to any of several other vertices to which it is connected by communication lines. Their algorithms for broadcasting can be described briefly as follows:

(i) transmission of separately addressed packets - this is the simplest and perhaps most often used, broadcasting scheme, whereby one copy of the message is routed to each of the n-1 other vertices;

(ii) multi-destination addressing - in this case fewer total messages are sent, but each of them is routed to a subset rise to their term "center based broadcasting". The reader is referred to the thesis of Wall [54] for an interesting discussion of several instances where the need to broadcast arises in the use of computer networks.

Finally, Santoro (1980) [46] and Korach, Rotem and Santoro (1980) [36] have designed a variety of algorithms for distributed networks in which, in effect, some form of broadcasting and/or gossiping is necessary in order to determine current topological information about the network. For example, by carrying out some form of broadcasting from each vertex, the center, the median, the radius, the diameter and a spanning tree of a network can be determined; inspite of the fact that each vertex has only local information about the global topology of the network.

7. Future studies

Although a moderate amount of work has been done to date on gossiping and broadcasting in communication networks, a substantial amount of work remains to be done, not only in extending the results in each of the existing areas, but in exploring a number of new areas. For example, the great majority of work-to-date has adopted the constraint that a vertex can only place one call per unit of time. Although a few papers have made the assumption that a vertex can simultaneously call all of its neighbors (shout) (eg. [11], [13], [26] and [49]), this variant is essentially unstudied.

Virtually all of the work to date has been done on local broadcasting and local gossiping, i.e. a vertex can only call a neighbor. Although Farley [15] showed that long distance broadvertices (one-to-many instead of one-to-all).

In conclusion, several things seem fairly clear relative to this survey. First, the subjects of broadcasting, gossiping and related information dissemination processes are terribly rich in real-world applications. Second, as we enter the age of the Information Society in which more than 70% of the work force will be occupied in one way or another with information processing, increasing degrees of sophistication will be required in our methods of information dissemination. Third, the mathematics of information dissemination has not yet been well developed, but it appears that it will involve a blend of combinatorics, graph theory, probability and computing. And finally, as the results in this survey suggest, our understanding of this subject at present is at best primitive.

8. References

[1] B. Baker and R. Shostak, Gossips and telephones, Discrete Math., 2(1972), 191-193.

[2] A. Bavelas, Communication patterns in task-oriented groups, J. Acoust. Soc. Amer., 22(1950), 725-730.

[3] G. Berman, The gossip problem, Discrete Math., 4(1973), 91.

[4] J. C. Bermond, Le problem des "ouvroires" (hypergraph gossip problem), Colloq. CNRS. Problemes Combinatoires et Theoire des Graphes, Orsay, France, 1976, 31-34.

[5] J. C. Bermond, private communication, January, 1981.

[6] R. Bumby, A problem with telephones, SIAM J. Algebraic Discrete Methods, to appear.

[19] A. Farley, S. Hedetniemi, S. Mitchell and A. Proskurowski, Minimum broadcast graphs, Discrete Math., 25(1979), 189-193.

[20] A. Farley, S. M. Hedetniemi, S. T. Hedetniemi and A. Proskurowski, Cumulative time gossiping, University of Oregon Tech. Rept. CIS-TR-81-7, August, 1981.

[21] A. Farley and A. Proskurowski, Gossiping in grid graphs, J. Combin. Inform. Systems Sci., 5(1980), 213-224.

[22] A. Farley and A. Proskurowski, Broadcasting in trees with multiple originators, SIAM J. Algebraic Discrete Methods, to appear.

[23] A. Farley and A. Proskurowski, Networks immune to isolated line failures, University of Oregon Tech. Rept., CIS-TR-80-19, 1980.

[24] A. Farley and A. Proskurowski, Extremal graphs with no disconnecting independent vertex set or matching, University of Oregon Tech. Rept., CIS-TR-80-20, 1980.

[25] M. R. Garey and D. S. Johnson, Computers and Intractability: a guide to the theory of NP-completeness, W. H. Freeman and Company, San Francisco, 1979.

[26] W. M. Gentleman, Some complexity results for matrix computations, J. Assoc. Comput. Mach., 25(1978), 112-115.

[27] M. Golumbic, The general gossip problem, IBM Research Report RC 4977, August, 1974.

[28] A. Hajnal, E. C. Milner and E. Szemerédi, A cure for the telephone disease, Canad. Math. Bull., 15(1972), 447-450.

[29] F. Harary and A. Schwenk, The communication problem on graphs and digraphs, J. Franklin Inst., 297(1974), 491-495.

[30] S. M. Hedetniemi and S. T. Hedetniemi, Broadcasting by decomposing trees into paths of bounded length, University of Oregon Tech. Rept. CS-TR-79-16, 1979.

[31] D. S. Johnson and M. R. Garey, private communication, May, 1978.

[45] A. Proskurowski, Minimum broadcast trees, IEEE Trans. Comput., C-30(1981), 363-366.

[46] N. Santoro, Determining topology information in distributed networks, Proc. Eleventh SE Conf. on Combinatorics, Graph Theory and Computing, Utilitas Mathematica, Winnipeg, 1980, 869-878.

[47] P. Schmitt, Spreading information by conferences, Discrete Math., 15(1976), 305-306.

[48] P. J. Slater, E. J. Cockayne and S. T. Hedetniemi, Information dissemination in trees, SIAM J. Comput., to appear.

[49] Q. F. Stout, Information spread in mesh-connected computers, presented at the SIAM Conf. on the Applications of Discrete Mathematics, Troy, New York, June, 1981.

[50] R. Tijdeman, On a telephone problem, Nieuw Arch. Wisk., 19 (1971), 188-192.

[51] F. L. Van Scoy, Parallel algorithms in cellular spaces, Ph.D. Thesis, University of Virginia, 1976.

[52] F. L. Van Scoy, Broadcasting a small number of messages in a square grid graph, Proc. Seventeenth Allerton Conf. on Communication, Control and Computing, 1979.

[53] J. Vuillemin, A data structure for manipulating priority queues, Comm. ACM, 21(1978), 309-315.

[54] D. W. Wall, Mechanisms for broadcast and selective broadcast, Stanford University, Computer Systems Laboratory Tech. Rept. 190, June, 1980.

[55] D. W. Wall and S. S. Owicki, Center-based broadcasting, Stanford University, Computer Systems Laboratory, Tech. Rept. 189, June, 1980.

[56] D. B. West, A class of solutions to the gossip problem, Stanford University Tech. Rept. STAN-CS-78-693, November, 1978.